

Telecommunications Systems

Background of the Invention

The present invention relates to telecommunications systems such as telephone networks comprising a plurality of interconnected telephone exchanges and communication therebetween.

Conventional telephone networks are fully-meshed in that each trunk exchange has a direct connection to every other trunk exchange. Traffic levels in telephone networks are increasing, leading to a continued need to increase the capacity of such networks. To achieve this, exchanges need to be enlarged and new exchanges installed. In fully-meshed networks, a significant number of additional network interconnections are needed with each new exchange, leading to further increases in the number of ports needed per exchange. This results in less efficient networks with costs per unit of traffic increasing.

15

Summary of the Invention

By using the present invention it is possible to provide a more efficient and more easily configurable network.

20

The present invention provides a telecommunications system comprising one or more nodes and a plurality of telephone exchanges wherein two or more of the telephone exchanges are arranged to communicate with each other via the one or more nodes; wherein communication via the one

653074-6 TEE 200

or more nodes is in the form of packets; wherein the one or more nodes comprise routers.

In a further embodiment the present invention provides a telecommunications system wherein communication via the one or more routers uses internet protocol (IP).

5

The present invention also provides an adapter for providing the above telephone exchanges with a means of inter-communication of traffic via the one or more nodes wherein the adapter comprises means for converting the traffic between packetised and non-packetised form.

10 Brief Description of the Drawings

Embodiments of the invention will now be described by way of example with reference to the drawings in which:

15 Figure 1 shows in block diagram form a typical, conventional, large telecommunications network of the prior art;

Figure 2 shows the network of Figure 1 with routers according to a first embodiment of the invention;

Figure 3 shows the network of Figure 1 with routers according to a second embodiment of the 20 invention;

Figure 4 shows in block diagram form a typical, conventional, small telecommunications network of the prior art;

Figure 5 shows the network of Figure 4 with routers according to a third embodiment of the

CONTINUATION

invention;

Figure 6 shows the network of Figure 1 partially upgraded with routers according to the first embodiment of the invention;

Figure 7 shows in block diagram form an adapter according to the present invention;

5 Figure 8 shows a possible internal organisation of the adapter of Figure 7.

Detailed Description of the Preferred Embodiments

Turning to Figure 1 a typical, conventional, large telephone network comprises a number of local exchanges L interconnected via trunk exchanges T. Each local exchange L is connected to two (or more) trunk exchanges T, which are fully-meshed. In current networks this fully-meshed trunk interconnect typically uses mainly PDH transmission.

Figure 2 shows the typical, large network of Figure 1 after trunk upgrade according to a first embodiment of the invention. In this embodiment two routers R are shown, although the number may vary in practice. Each router R is connected to every trunk exchange T. Rather than a large number of low bandwidth connections, the trunk interconnect now consists of a few high bandwidth connections.

20 Figure 3 shows the typical, large network after local upgrade. With the addition of some simple multiplexers M, all the local exchanges are directly connected to the routers R. The trunk exchanges (not shown) may still be used to provide regional interconnect and interconnection to non-upgraded exchanges and other networks.

Figure 4 shows a typical smaller network. The smaller network has tandems N rather than trunk exchanges, which may have a few customers directly connected. Each local exchange L is connected to several (possibly all) tandems N, there is no interconnection of the tandems.

- 5 Figure 5 shows the typical, smaller network of Figure 4 after being upgraded according to a third embodiment of the invention. Routers R are installed, similar to those used to upgrade the large network, interconnecting the local exchanges L, through multiplexers M. The old tandems can then be reused, for example, as local exchanges and points of interconnect to other networks.
- 10 The routers R can be installed as a trunk network replacement piecemeal. Figure 6 shows the large network of Figure 1 partially upgraded. Here traffic is carried by a mixture of the old trunk interconnect and new routers R. A number of trunk exchanges T have connection to the routers R and are able to communicate with others so connected via the routers. However not all trunk exchanges T are connected to the routers R at this stage and the old trunk interconnect
- 15 is still used for communication between exchanges not connected to the routers R and for communication between each such exchange and each of the exchanges connected to the routers R. The number of routers R can be upgraded from one to four or more as further exchanges are connected.
- 20 In this way the routers may be extended to eventually connect all existing trunk exchanges.

It may be desirable to move some of the existing transmission over to the router network, to take advantage of the greater efficiency of the routers R compared with the existing meshed

PCT/GB2007/001250

PDH, and to overcome a shortage of physical ports on the conventional switches.

Each exchange requires routing data containing information on the way the exchanges are interconnected within the telecommunications system. Every time an exchange gets connected

5 to the routers R, all the other exchanges that have already been connected to the routers will need to have new routes set up for communication with the newly connected exchange. Re-configuring exchanges in this way every time another exchange is connected via the routers is expensive. A preferred solution is for routing data relating to all exchanges to be loaded into each exchange as it is connected to the routers, but only to enable data relating to routes to those
10 exchanges that have already been so connected.

Hence the implementation of the present invention simplifies the data requirements in exchanges. On installation, every exchange (trunk or local) can have exactly the same network data loaded. This data can be pre-prepared for the whole network with the data for the
15 exchanges connected to the routers being enabled as the network is enhanced. This data includes the network part of the digit decode, which is needed to determine the exchange to handle the call, and the route to be used. The same route number can be used for all exchanges to get to the same far-end exchange. The digit decode requirements of the exchanges would grow to support routing to all local exchanges.

20

The existing trunk exchanges, could still be used, e.g. as regional trunk exchanges (DJSUs), allowing reduction or elimination of sideways routes between local exchanges. If at any time these regional trunk exchanges are overloaded, the traffic can be handled in the main trunk

Subj b

BB
cont
(router) network.

Extension of the routers to directly interconnect local exchanges may be implemented once there is a sufficiently developed provision of trunk routers to provide the interworking.

5

When expanded to connect the local exchanges, enhancement may be necessary to carry the signalling directly over the IP connections. One way of interworking to local exchanges and other networks is to have simple interworking routers, which are effectively minimalist exchanges. These could be the existing trunk exchanges.

10

The adapters may use ATM, IP, or virtually any packet-based network. In many large-scale networks, IP is carried over ATM and a similar scheme could be used in the network of the invention. Alternatively IP could be carried over point-to-point protocol (PPP) over SDH or SONET, or any other appropriate interface. In the network of the invention, data is sent across 15 the network to an address determined by the adapter. No call processing is necessary within the ATM, IP or equivalent packet networks. In effect the packet (e.g. ATM or IP) network forms a transport multiplex layer below that of the individual PSTN circuits. The packet (e.g. ATM or IP) network transport multiplex provides an alternative to the rigid PDH transmission network that is more flexible in the way it multiplexes channels. Conceptually the 20 interconnected adapters may be thought of as providing an intermediate layer between the PSTN circuits on top and the transport multiplex underneath.

When operating over an IP network, the adapter could use the same AAL2 multiplex structure

that it would use over ATM or it could use similar multiplex structures to those used by VoIP.

AAL2 has been designed to carry low rate, short length packets in delay sensitive applications in a bandwidth efficient way. The AAL2 multiplexes many separate low speed streams such 5 as voice, data and signalling together on a single IP packet. AAL2 has been standardised in ITU-T I.363.2. AAL2 is an efficient way of multiplexing low-speed circuits together and is suitable for use over IP in the network of the present invention.

Voice traffic is given a high priority. This ensures that the delay variation would be well constrained, thus limiting the amount of additional buffering required to remove the effect of cell delay variation at the receiver.

The use of VoIP over ATM to provide the interconnection of adapters over the PTO network essentially uses a subset of the general VoIP transport as performed by any conventional VoIP gateway.

The proposed solution supports carrying voice as AAL1 or AAL2 in configurations to suit different networks at different stages of evolution.

20 Voice may be carried across the network of the present invention in more than one way. Where the adapters communicate using IP native-mode, voice may be packetised in a similar way to AAL2 for sending over the IP network, or packetised as voice over IP (VoIP). Voice over AAL2 may be implemented using the ITU-T I.366.2 "AAL type 2 service specific convergence

sub-layer for trunking" speech multiplexing and encoding principles. This is efficient and supports transition to low bit-rate coding schemes, transcoding, etc.

Fundamentally the adapter working over IP (IP adapter) functions in exactly the same way as

5 the adapter working over ATM (ATM adapter). As there is no call set up mechanism in the IP network, an adapter sending traffic over IP has to assume that the transmission capability it needs is always available. As with the ATM adapter, there is a need to ensure that the IP voice traffic is given a high enough priority to enable an adequate quality of service to be achieved. In order not to adversely affect the quality of service seen by the PSTN network, the IP network 10 delay for voice traffic must be relatively low. This means that there must be a means to effectively prioritise voice traffic. The adapter needs to process voice traffic with a higher priority than data traffic in order to prevent unacceptable delay to voice traffic.

This may be achieved in a number of ways, for example if IP is carried via ATM where 15 prioritising of voice traffic is supported. Operation over a private (totally-dedicated network) where voice traffic does not have to share the network with data traffic allows an adequate quality of service to be easily achieved by keeping delays to an acceptable level.

The invention may be implemented over IP native-mode which is capable of carrying some non- 20 voice traffic but care is necessary to ensure that it does not affect the low delay requirements of the voice traffic.

In order to minimise delay when using IP, the length of IP packets carrying voice may be

DOCUMENT EDITION

shortened. However, as each packet contains its own header bytes, this is wasteful of the network bandwidth due to the increased proportion of header information to useful data sent. Advantageously a number of different circuits may be multiplexed into a single, large packet. In this way delay is still reduced without increasing the proportion of header information to data sent.

Use of AAL2 carrying a number of different channels multiplexed together can provide low delay (as in the case of short packets) without excessive amounts of header information.

10 A single network according to the present invention could use ATM to provide traffic to a first destination whilst at the same time using IP over ATM to provide traffic to a different destination. Hence PSTN circuits could be carried across the PTO network using either ATM or IP over ATM on a route by route basis. In doing this a relationship is set up between individual PSTN circuits and routes across the PTO network.

15

Telephone calls include traditional hand-set to hand-set calls but also calls initiated or answered by machine including computers. Such calls may contain voice and/or data from e.g. modems or facsimile equipment. Handling of such calls include call set-up, routing and clear-down. No call handling is required in the central routers. All the call handling is handled in the PSTN,

20 using existing protocols with all the existing features.

To achieve network resilience in the presence of equipment failure, it is desirable to have at least two, and preferably more central nodes so that if one node fails the network would still be able

DOCUMENT EDITION 00

to work via the other nodes. Here the term "node" is used to indicate a cross-connect, router, or other suitable packet-based switch. In an IP based network this would require a plurality of truly independent IP routers which were geographically separated and independently controlled and managed. They would also need to be resilient against unauthorised intrusion and hacking.

5

Between each pair of exchanges in an IP based PTO network, a Logical Route would be established, that corresponds to all the active calls between that pair of exchanges. This Logical Route is dynamic in nature and the bandwidth of the Logical Route changes to match the changing demand as existing calls terminate and new calls start.

10

In a preferred embodiment of the present invention, each trunk/local exchange has an STM1 (155M) connection to each router, carrying one Logical Route per destination exchange. For most local exchanges these 155M connections can be shared between several exchanges, as detailed below. The Logical Route can grow or shrink as necessary: the total quantity of traffic 15 is only limited by the capacity of the transmission to carry it. As traffic patterns change, Logical Routes may grow or shrink on a dynamic basis: no configuration is necessary.

The exchanges communicate using existing SS7 protocols, however there is an option (that is necessary when it is extended to local exchanges) to carry the MTP layer 3 over the Logical

20 Route that connects each exchange, thus eliminating the need for any other signalling hardware.

Further bandwidth savings may be achieved if only active PSTN circuits (i.e. those carrying information) are packetised. This requires a busy/idle timeslot map to be carried in an end to

6620077-1675224-160

end control channel using spare capacity within the IP Logical Routes. This technique allows the mapping to be changed dynamically and in a similar way to that set out in detail in patent number GB 2276518 in the name of GPT Limited. Optional echo cancellation may also be provided.

5

Optional signalling transfer point (STP) function to handle SS7 signalling may be provided. This feature will extract redundant data (Level 2 interframe fill) from the SS7 signalling stream.

Optional modem (fax) detection may be provided. Low speed modems are not bandwidth 10 efficient. In particular, many connections, e.g. fax calls, are essentially unidirectional. This optional feature will detect such calls and convert the modem signals to baseband data for packetisation into IP, taking advantage of the fact that communication across IP can be asymmetric. Thus further bandwidth savings can be achieved.

15 Although this has been described in relation to improved trunk networks, the present invention could also be deployed into many other fields for both public and private networks.

Similarly the present invention has application to leased lines within public networks. Leased lines are often under-utilised, resulting in idle timeslots being carried across the network. The 20 use of the present invention, will allow the network operator to only transport the active channels without degrading the service presented to the end user.

The present invention provides for straightforward enhancement to support traffic growth

GPT-61683-USP

wherein the majority of the routing data needed by the PTO exchanges is identical. This allows operators to add to and evolve current networks as and when enhancements to performance or capacity are required.

5 For interfacing with sources of non-packetised traffic an adapter is provided, see Figure 7. The adapter functions to convert traffic for sending to the routers into packets and to re-convert packets received from the routers to un-packetised form. In one embodiment the adapter provides interworking between a PSTN trunk (G.703/G.704) and an IP domain to allow IP to be used to provide an efficient and flexible interconnect between exchanges by using variable
10 bandwidth Logical Routes.

For interfacing with STM networks, the adapter provides interworking between IP and STM, with preferably one adapter for each route via the routers. These adapters perform interworking between 2M bit/s PCM traffic and IP and also provide a means to connect signalling directly
15 with each of the other local exchanges. These adapters communicate with each of the processor clusters as most of the signalling traffic passes through them.

We now consider the adapter construction in more detail. Referring to Figure 8 by way of example, each adapter unit comprises a housing with a PSU, a common card and up to 4 PCM
20 cards, although more could be added. Each PCM card could handle up to 16 PCMs, the common card has an STM1 interface, an Ethernet interface and a controller for the PCM cards. Each exchange could typically need 4 adapters in this embodiment and these could preferably be mounted two per shelf of the exchange.

00000000000000000000000000000000

The routers operate under PSTN control and provide no control processing themselves. Advantageously a small number of cross links are provided between the main routers, to provide some flexibility after major transmission network outages. These would only be brought into use when a pair of local exchanges have no common accessible trunk routers (a very rare 5 occurrence).

To control the routing and use of circuits, the exchange's call processing system (CPS), sends mapping information to the adapter function at the start and end of each call. At the start of each call, the mapping message gives the association between the PCM circuit and the PSTN 10 side with the route and circuit on the logical side. For ATM based PTO the route would be a VP or a VC, for IP based PTO the route would be the Logical Route to the destination exchange, the mapping message may also indicate the type of call, so that appropriate compression can be used. At the end of the call the CPS function will send another message to cancel the mapping function.

15

A processor interconnect handles the signalling traffic, the mapping messages from CPS and general management. A suitable interface to use here is Ethernet as it allows the interface to be fairly open for use on other platforms if necessary.

20 The hardware of the PCM termination could support both E1 and T1 or be dedicated to either data rate.

The adapters also exchange status information with each other to ensure that each is operational

and to exchange and pass on traffic-blocking information. Periodically and when necessary (for example because of a status change) each adapter should send a cell to each far end adapter.

This cell contains three elements:

- 5 1. The adapter receiver status (i.e. whether the adapter is happy with what it is receiving from the far end adapters);
- 10 2. A hold indication that, if set, should cause the far end adapter to stop the flow of new traffic to that adapter;
3. Status information on the PCMs to the exchange (i.e. reflecting which PCM/circuits can be used).

When an adapter receives a check cell, it reacts as follows:

Sub B3

- 15 1. If the far is not happy, it should force clear all calls and hold that route busy (but still send its own check cells);
2. If the far end requests a hold, send a hold to the associated exchange;
- 20 3. If the far end ceases to request a hold clear the condition (the adapter may delay reacting to a change of state to prevent the network oscillating);
4. If the far end has PCM failures, reflect these on to the alarm indication signal (AIS)

of its transmitted PCMs.

The Control protocol between the call processing and the adapter units is implemented as an application programming interface (API). This API covers the setting up and breaking down 5 of switching associations, and optionally management (e.g. failure and configuration management).

Management of the adapter includes managing the equipment, configuring it and handling any appropriate statistics. Management is preferably achieved through the exchange or, when the 10 adapter is being used as a standalone device, a Simple Network Management Protocol (SNMP) interface would be appropriate for it to be managed by third party management systems. Alternatively a standalone adapter could be managed through the SDH, for example by using the Equipment Management Operations System (EMOS) particularly when the adapter is closely associated with transmission equipment. Thus, depending on where and how the adapter 15 is being used, different management solutions are necessary.

When the adapter is being used as standalone equipment, a SNMP management interface over the Ethernet control port will be appropriate. This will be controlled by a Managed Information Base (MIB) which should be published. The data model necessary for these applications may 20 differ significantly depending upon the type of exchange used. Flexibility is essential to allow for the addition of functionality such as transcoding, packing algorithms and data rate change. When the adapter is deployed along with SDH transmission, it may be appropriate to manage it through EMOS and the rest of the SDH. This could be through an Ethernet port, through the

662077-072200

DCC of the STM1, or through the IP.

In principle all the traffic could be handled by a single router, but this is likely to be unacceptable in practice for availability reasons.

5

There will preferably be at least two routers to allow for catastrophic failure. However even with two it will be necessary to have a very high level of redundancy with alternative routing for the transmission to these routers. With four routers there is a good level of built-in network redundancy, without the cost being too high. Each should preferably be dimensioned such that

10 any two could handle the normal peak traffic loads of the entire network.

For maximum availability suitably diversely routed transmission should be arranged to the routers with each router on an independent site.

15 The basic architecture would work with any of the following combinations:

Two routers connected over two diverse routes and using transmission protection; three routers over two diverse routes using transmission protection or three routers with three diverse routes without transmission protection. Four routers over two diverse routes without using 20 transmission protection; four routers over two diverse routes using transmission protection and six routers connected via two to three diverse routes without transmission protection. Obviously the numbers above could be varied by the skilled worker depending on the requirements concerning resilience to failure and cost.

Each adapter would need at least one IP address for routing traffic across the network. Each adapter may have additional addresses for other purposes.

5 The arrangement described above, using IP, can bring benefits in the form of transmission saving, a dynamic network, data simplification and switch port saving.

Transmission saving results from the use of large units (e.g STM1 rather than 2Mbit/s such as VC12) which only has to be routed to a far-end adapter via one of the central nodes (e.g. ATM 10 cross-connect or IP router). In many cases these large transmission units may be part empty thus allowing lower priority data services to be carried. This results in a much more efficient use of transmission than in traditional, directly-connected networks.

By dynamic network is meant that traffic patterns can change on a rapid (e.g. minute-by - 15 minute) basis without the need for complex load balancing or for guessing future traffic loads on particular routes. This renders the network efficient and easy to manage.

Data simplification results from being able to use the same routing and network data in all exchanges. This data is also simpler, reflecting the simpler structure of the PTO network.

20

The number of switch ports required in exchanges can be reduced due to the removal of the physical mesh interconnection.

062007-072009